

# An Interactive Graphics Display Architecture.

Matthew Regan and Ronald Pose.

Department of Computer Science, Monash University.  
Clayton, Victoria, Australia 3168.

E-mail:	regan@bruce.cs.monash.edu.au	rdp@bruce.cs.monash.edu.au
Phone:	+61 3 565 5812	+61 3 565 5203
Fax:	+61 3 565 5146	+61 3 565 5146

## Abstract.

In recent times Virtual Reality has become a growing area within computer graphics with many and varied applications. Many systems suffer from inadequate performance resulting in a lack of realism and even psychological side effects[1]. A large body of research exists to improve the performance of Virtual Reality systems by increasing the rate at which scenes are drawn. This project differs in that the display controller is redesigned so as to cater directly for the needs of a Virtual Reality display system.

By performing complex computation within the display controller as a pixel is being displayed, latency to user head rotations may be directly reduced and latency to translations and stereoscopy may be indirectly reduced. Simulations have shown this architecture to be feasible to implement with current off-the-shelf technology. A prototype is currently being constructed[2].

## 1.0 Introduction.

This paper introduces and briefly investigates a new computer graphics display controller architecture which is suited to the requirements of head mounted Virtual Reality (VR) display systems. This architecture differs from standard display controllers in the nature of the video memory addressing mechanism. Rather than fetching pixels for display from the video memory sequentially, the fetch mechanism bases the pixel fetch location on the pixel's actual screen location (as seen through the wide-angle viewing lenses) and the orientation of the user's head. This means the viewport mapping is delayed until after the scene has been rendered.

The new display controller directly reduces latency to user head rotations by fixing user head rotation latency to the refresh period of the display device. Latency to other forms of interaction such as user translations through the virtual scene, head motion parallax and stereoscopy are indirectly reduced through the use of image overlaying.

Image overlaying or image composition [3] is a technique often used by low-cost non-head mounted display graphics systems such as video games to increase the apparent realness of a scene. It is also often used by high performance systems to increase the maximum rendering rate. Different sections of the visible scene are drawn into separate display memories then overlayed to form a final scene (a similar technique is often used when creating cartoons to reduce the drawing required per display frame). When translating through the virtual scene, each display memory may be updated independently. Close objects receive more of the rendering engine's time than less frequently changing background objects. Image overlaying generally does not work very well with conventional display controllers in VR systems as user head rotations invalidate all display memories including those containing background images. This mass invalidation occurs because all of the images in the display memories are rendered with a fixed viewport mapping so when the viewport changes due to a user head rotation all of the images have an incorrect viewport mapping. The new display controller does not suffer the same

drawback since the rendered scenes in display memory are independent of the viewport mapping orientation and do not become invalid when the viewport mapping changes<sup>1</sup>.

## 2.0 A New Display Paradigm.

Display generation usually consists of two processes. The first process called rendering draws images into display memory. The second process sends the image in the display memory to the display device. This process is performed by the display controller. Such a display system is shown in Figure 2.1.

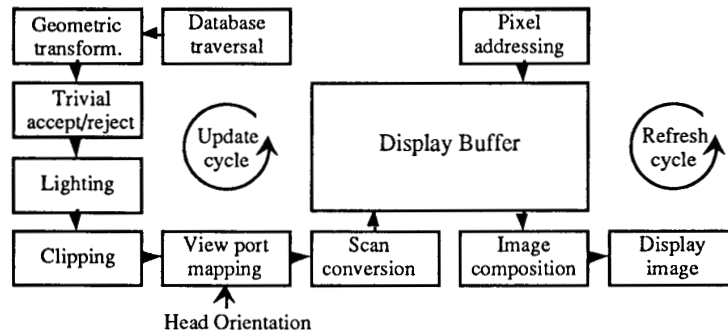


Figure 2.1 A typical display system.

Figure 2.1 clearly shows two cycles, the update cycle and the refresh cycle. An update cycle has occurred when all required primitives in a scene have been drawn into display memory, where as a refresh cycle has occurred when an image in the display memory has been sent to the display device. Conventional VR display systems bind the latency to user head rotations to the period of the update cycle. While this has the obvious advantage of not requiring any specialised hardware, latency tends to suffer as the update cycle period is usually lengthy, unpredictable and dependent on scene complexity.

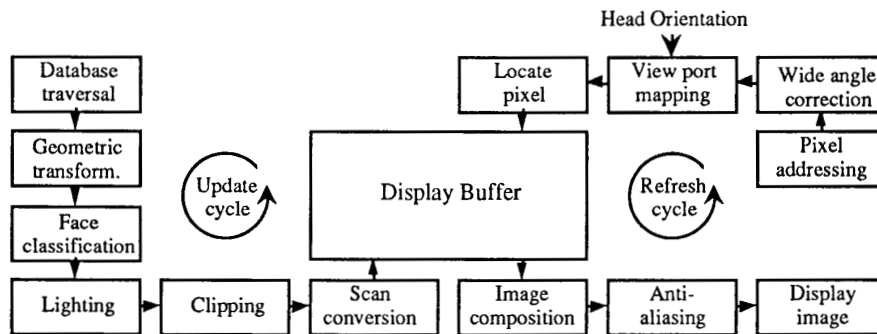


Figure 2.2 A model of the new display system.

If we modify the architecture of the display controller it becomes possible to bind rotational latency to the refresh cycle period which tends to be short, fixed in length and independent of scene complexity. In Figure 2.1 the user's head orientation was introduced in the update cycle,

<sup>1</sup>The only exception to this is when providing stereoscopy for large orientation changes. Some re-rendering of close objects may be required for an accurate stereoscopic view of the scene.

while Figure 2.2 depicts a display architecture in which the latency to user head orientation is bound to the refresh cycle period.

Although many projects focus on improving latency by reducing the period of the update cycle [4], or providing faster access to the display memory [5], few projects alter the architecture of the display controller.

Viewport mapping is performed after rendering so the image in display memory must be as independent of user head orientation as possible. This means the image must completely encapsulate the user's head. Within this architecture, the surface of a cube centred at the user's head was the chosen topology for storing images within the display memory. Figure 2.3 shows an image drawn onto the surface of a cube as it is stored in display memory as well as view from the image at an arbitrary orientation. Note that the distortion of the lower right hand side of the pool table in the display memory image is not seen in the view created from the image.




Figure 2.3 An image stored in display memory and an arbitrary view from the image.

### 3.0 The Address Recalculation Pipeline (ARP).

Traditionally the display controller performs relatively simple computations, however the operations performed by the controller may be made more complex provided the computation may be pipelined. For current generation medium resolution displays the pipeline hardware on the controller must have a throughput of one pixel every 40 nanoseconds.

The purpose of the address recalculation pipeline is to provide non-sequential video memory addressing mechanisms which automatically correct for wide-angle viewing lenses and user head orientation. The pipeline orientates the address generation process within display controller with the orientation of the user's head as the pixels are being fetched from video memory. Due to the nature of the computation, the pipeline has been dubbed the Address Recalculation Pipeline (ARP). A block diagram of the ARP for the left display device is given in Figure 3.1.

The first input to the pipeline is the screen location of the current pixel being fetched. A look-up table converts this screen location into a three dimensional cartesian vector pointing in the direction at which the pixel is seen relative to the user's head. This vector is then multiplied by a matrix containing user head orientation information which results in another three dimensional vector pointing in the direction at which the pixel is seen relative to the world co-ordinate system. The vector is then converted into a two dimensional video memory location. The pixels at the resultant memory location are fetched from the multiple display memories and composed before being anti-aliased. Finally the RGB pixel to be displayed is produced. Both the left and right displays have an address recalculation pipeline.

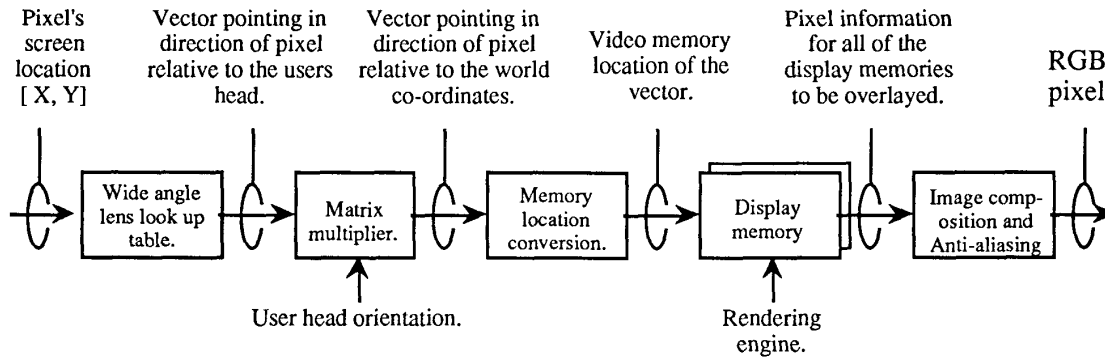


Figure 3.1 An Address Recalculation Pipeline.

### 3.1 Wide Angle Viewing Lenses Look-up Table.

The first stage of the ARP must convert a pixel's screen location into a unit cartesian vector pointing in the direction at which the pixel is seen by the user relative to the head mounted display. Due to the complexity of such a transformation (especially when unusual wide-angle lenses are used [6,7]) a look-up table is used for this stage. Many wide-angle viewing lenses preserve a standard viewport mapping, however if a special mapping is required for higher fields of view, the look-up table may be loaded with a new lens mapping and the lens distortion may be compensated for without run-time penalty. Figure 3.2 shows views created from the display memory image in Figure 2.3 with different wide-angle viewing lens mappings.




Figure 3.2 Lens distortion mappings in the wide-angle viewing lens look-up table.

### 3.2 Matrix Multiply Stage.

The matrix multiplier stage consists of two matrix multipliers (one per eye) that are used to transform the user's head orientation to the world orientation of the display memory.

The unit vector  $[X_i \ Y_i \ Z_i]^T$  in Figure 3.3 points in the direction at which the pixel is seen by the user relative to the head mounted display. The matrix on the left hand side of the multiplication in Figure 3.3 contains three orthogonal unit vectors orientated in the direction of the user's head relative to the world coordinate system and is supplied directly from the user head tracking equipment at the start of each refresh frame. Once the orientation matrix is multiplied by the pixel direction vector a resulting output vector is produced. This output vector

$[X_0 \ Y_0 \ Z_0]^T$  is a unit vector pointing in the direction at which the pixel is seen by the user relative to the world coordinate system.

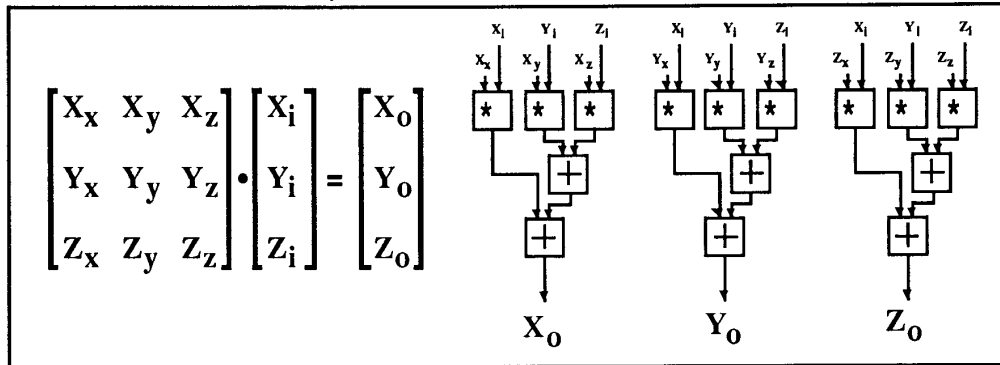


Figure 3.3 One of the matrix multipliers.

### 3.3 Display Memory.

In order to understand how the vector conversion stage of the pipeline operates the topology of the display memory must be known. A requirement of the display memory is that it must completely encapsulate the user's head. The first most apparent solution is to use a spherical encapsulation surface with unit radius where lines of longitude and latitude are mapped to the rows and columns of the display memory. While this approach works it has some drawbacks. First, expensive Digital Signal Processing devices are required for vector conversion, secondly pixel sizes vary greatly depending upon the pixel's position on the surface of the sphere, resulting in inefficient usage of video memory.

The topology chosen to encapsulate the user's head within this architecture is the surface of a cuboid with side length of two units, centred at the origin. This means the display memory is divided into six faces where each face contains a two dimensional grid of pixels. When folded correctly these faces form a complete surface around the user's head.

A run-time advantage of selecting the surface of a cube over the surface of a sphere is that the scan conversion process for the cube requires minor modifications to existing scan conversion/shading algorithms while scan conversion for the spherical surface requires major, computationally expensive changes to common algorithms.

### 3.4 Vector Conversion stage.

The vector conversion stage within the pipeline must convert a 3D unit vector into a display memory location. This conversion involves computing the point at which the ray intersects with the surface of a cube. This conversion involves computing the point at which the ray intersects with the surface of a cube. Aligning the cube to the axis of the co-ordinate system such that each face of the cube has one of its X, Y or Z co-ordinates fixed at  $\pm 1$  means that the point of intersection may be computed with a division followed by a set of range checks. For example if  $y/x$  and  $z/x$  are both in the range  $(-1.0 \ 1.0)$  the ray may intersect with only two of the six faces. The sign of  $x$  then determines which of the two faces the ray intersects, with the point of intersection being  $(y/x, z/x)$  on the face.

Finding dividers which operate at video speeds is quite difficult and the process is further complicated by the fixed point nature of the pipeline. Consequently a table look up followed by a multiplication was the division method chosen for this architecture. Sixteen bit precision inputs/outputs require 128 Kbyte look up tables. Using a look-up table means that the range of outputs may be bounded between  $\pm[1.0 \text{ and } 2.0]$ , simplifying preliminary face classification (a

candidate face must have a divisor in the range  $\pm[1.0 \text{ to } \sqrt{3}]$ . Adders are used post division to simplify signs for range checking.

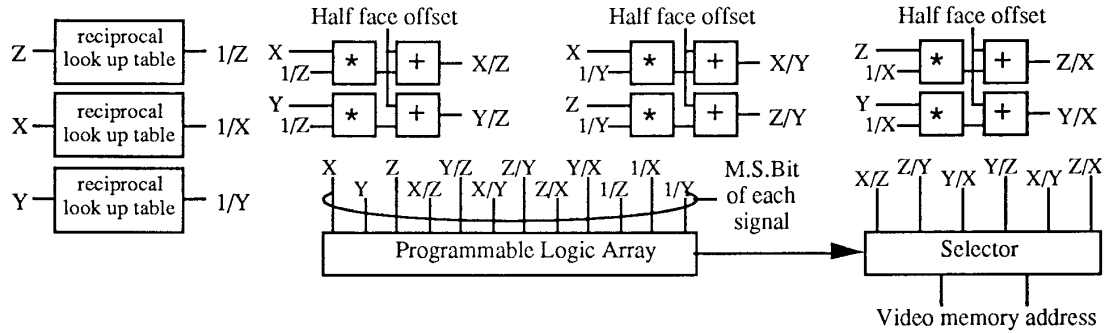


Figure 3.4 Block diagram of one of the vector converters (some pipeline delay registers have been removed for clarity)

A programmable logic array is used to classify face intersection from range information from all of the appropriate data paths. Aligned fixed point arithmetic with bounded reciprocal look up tables means that all required range information is stored within the top bit of each data path. Such a vector converter is depicted in Figure 3.4

### 3.5 Image Composition and Anti-Aliasing.

Image composition [3] is a technique whereby multiple display memories are fetched simultaneously with the same pixel location and the pixel with the smallest value in its Z-buffer is displayed. The Z-value comparison is usually performed by a tree comparator. In this application, different objects may be rendered into different display memories. These objects may then be re-rendered independently as required. In some circumstances a background scene may require no re-rendering (in such a case the background may be rendered using a high quality rendering technique such as ray-tracing). In actual practice the performance improvement obtained by image composition depends on how evenly objects may be clustered by their required update rate for a finite number of display memories.

Anti-aliasing is quite important in this architecture due to edge effects caused by double sampling. So as well as the typical staircased effect for edges running at an angle, edges which are horizontal or vertical also break up. The breaking up of horizontal or vertical edges seem to be more noticeable than the staircasing of edges running at angle. The solution chosen for this architecture is to fetch a block of four adjacent pixels simultaneously then apply some smoothing filter. The nature of the filter is determined by the redundant bits in the computation of intersection. That is, the position of intersection is known to sub-pixel accuracy, so the sub-pixel information is used to apply a filter.

### 4.0 Performance.

A complete formal analysis of the performance of this architecture is well beyond the scope of this paper, however some figures may be extrapolated from our architecture simulator.

Latency to head rotations (except stereoscopy) is fixed to the display update rate. This means the eye may tracking an object as the head rotates with little detectable latency. A display update rate of 60 Hz combined with a head tracking update rate of 120 Hz results in an average rotational latency of approximately 20 ms. Forward prediction of head location using a linear or higher order models [8] may reduce this latency even further. As the viewport mapping process does not have to pass through the double buffer display switching (as it independent of

rendering), the time at which the raster passes through the centre of an image is known precisely. This in turn improves the quality of the head predication process.

Providing a binocular stereoscopic view of the scene may require some objects to be re-rendered. Head rotations cause stereoscopy to diminish at  $\pm 90$  degrees of the original orientation, however due to the nominal intraocular distance of 6.3 cm for humans, only relatively close objects translate significantly, so image composition allows these objects to be re-rendered independently. Added to the fact that the human visual system takes several tens of milliseconds to perceive a stereo scene, and apparently is not sensitive to small temporal errors [9] the latency involved in re-rendering for stereoscopy is important but not critical to system performance.

Latency to translations is often less noticeable than latency to rotations as many translations are the result of a user gesture such as pointing a finger or pushing an accelerator rather than actual physical motion. Providing an adequate update rate for translation however is important to provide a sense of fluid motion. We are currently examining many update models so different objects may be re-rendered to different display memories in order to reduce the overall load on the rendering engine. Most of the update models are generally based on the size, distance and relative velocity of a particular object. Clustering algorithms are being examined to tie objects with a similar update rate to the one display memory.

## 5.0 Conclusion.

This paper outlines a technique for redesigning a display controller within a traditional graphics display system. By performing complicated computations within the display controller the system may be optimised for the needs of a Virtual Reality display system. Wide angle viewing lens effects and response to user head rotations may be compensated for within the display controller, avoiding the rendering bottleneck. The architecture allows image overlaying to provide more realistic scenes and a mechanism for focussing the rendering engine's time on sections of the scene which change the most frequently.

## Acknowledgment.

Matthew Regan acknowledges the support of an Australian Postgraduate Research Award.

## References.

- [1] Deyo, R., and Ingebretson, D., "Notes on Real-Time Vehicle Simulation," in implementing and interacting with Real-Time Microworlds, Course Notes 29 for SIGGRAPH 89, Boston, MA, August 1989.
- [2] Regan, M. and Pose, R., "A Low Latency Virtual Reality Display System", Tech. Report No. 92/166, Department of Computer Science, Monash University, 1992.
- [3] Molnar, S., "Image Composition Architectures for Real-Time image Generation" University of North Carolina Ph.D.dissertation, 1991.
- [4] Fuchs, H., et al. "Pixel-Planes 5: A Heterogeneous Multiprocessor Graphics System Using Processor Enhanced Memories," SIGGRAPH 89, pp. 79-88.
- [5] Whitton, M. C., "Memory Design for Raster Graphics Displays." IEEE Computer Graphics and Applications, Vol. 4, No. 3, March 1984, pp. 48-65
- [6] Robinett, W., and Roland, J., "A Computational Model for the Stereoscopic Optics for Head Mounted Display". In Presence 1, (Winter 1992), pp. 45-62.
- [7] Deering, M., "High Resolution Virtual Reality", SIGGRAPH 92, pp. 195-202.
- [8] Friedmann, M., Starner, T., and Pentland, A., "Device Synchronisation Using an Optimal Linear Filter". In Proceedings of the ACM Symposium on Interactive 3D graphics (Cambridge, Massachusetts, March 29 - April 1, 1992), pp. 57-62.
- [9] Lipton, L. "Temporal Artifacts in Field-Sequential Stereoscopic Displays". In Proceedings of the SID 22 (May 1991), pp. 834-835.